

Adaptive nonparametric detection in cryo-electron microscopy

Langovoy, Mikhail

*Max Planck Institute for Intelligent Systems, Department Empirical Inference
and Max Planck Institute for Developmental Biology*

Spemannstrasse 38

D-72076 Tübingen, Germany

E-mail: langovoy@tuebingen.mpg.de

Habeck, Michael

*Max Planck Institute for Intelligent Systems, Department Empirical Inference
and Max Planck Institute for Developmental Biology*

Spemannstrasse 35

D-72076 Tübingen, Germany

E-mail: michael.habeck@tuebingen.mpg.de

Schölkopf, Bernhard

*Max Planck Institute for Intelligent Systems, Department Empirical Inference
Spemannstrasse 38*

D-72076 Tübingen, Germany

E-mail: bernhard.schoelkopf@tuebingen.mpg.de

1 Introduction

Cryo-electron microscopy (cryo-EM) is an emerging experimental method to characterize the structure of large biomolecular assemblies. Single particle cryo-EM records 2D images (so-called micrographs) of projections of the three-dimensional particle, which need to be processed to obtain the three-dimensional reconstruction. A crucial step in the reconstruction process is particle picking which involves detection of particles in noisy 2D micrographs with low signal-to-noise ratios of typically 1:10 or even lower. Typically, each picture contains a large number of particles, and particles have unknown irregular and nonconvex shapes.

2 Statistical model

Suppose we have a noisy two-dimensional pixelized image. In the present paper we are interested in detection of objects that have an *unknown* colour. This colour has to be different from the colour of the background. Mathematically, we formalize this as follows.

We have an $N \times N$ array of observations, i.e. we observe N^2 real numbers $\{Y_{ij}\}_{i,j=1}^N$. Denote the true value on the pixel (i, j) , $1 \leq i, j \leq N$, by Im_{ij} , and the corresponding noise by ε_{ij} . Suppose that the noise on the whole screen is i.i.d. with an *unknown* distribution function F . According to the above,

$$(1) \quad Y_{ij} = Im_{ij} + \varepsilon_{ij},$$

where $1 \leq i, j \leq N$, and $\{\varepsilon_{ij}\}$, $1 \leq i, j \leq N$ are i.i.d., and Im_{ij} denotes the true value on pixel (i, j) . Assume additionally that for the noise

$$(2) \quad \varepsilon_{ij} \sim F, \quad \mathbb{E} \varepsilon_{ij} = 0, \quad Var \varepsilon_{ij} = \sigma^2 < +\infty.$$

Our next assumption is inspired by images from cryo-electron microscopy. We assume that pixels in all particles have color intensity b , where $b > a$. In other words,

$$(3) \quad \begin{cases} Im_{ij} = b > a, & \text{if } (i, j) \text{ belongs to a particle;} \\ Im_{ij} = a, & \text{if } (i, j) \text{ does not belong to any particle.} \end{cases}$$

Here we assume that a and σ^2 are also both unknown. Estimating a is one of the problems that are solved in this paper. The particle intensity b is also not known, and one of our goals is to construct a good estimate of b as well. The difficulty is that a priori we do not know which pixels contain pure noise and which pixels actually belong to some particle. Moreover, locations, shapes and exact sizes of particles are assumed to be unknown. The number of particles is also unknown, and in this paper we make no probabilistic assumptions about this number or about the distribution of particle locations. Thus we consider the case of a fully nonparametric noise of unknown level.

We make now the following crucial assumption. Let there exist a square K_0 of size at least $\varphi_0(N) \times \varphi_0(N)$ pixels such that K_0 doesn't intersect with any particle, and $\lim_{N \rightarrow \infty} \varphi_0(N) = \infty$. Here we only require that somewhere, in between the particles, there is one single square of size at least $\varphi_0(N) \times \varphi_0(N)$, such that this square is completely filled with noise. There could be several squares of this kind, or there could be bigger squares filled with noise. Locations of noisy squares are unknown. At the first step, our spatial scan estimator is going to find at least one of those noisy patches.

If a pixel (i, j) is white in the original image, denote the corresponding probability distribution of Y_{ij} by P_0 . For a black pixel (i, j) we denote the corresponding distribution by P_1 . For any set of pixels K , we denote its cardinality by $|K|$. The empirical mean \bar{Y}_K of all the values in K is

$$(4) \quad \bar{Y}_K = \frac{1}{|K|} \sum_{(i,j) \in K} Y_{ij}.$$

The total sum of values in K is denoted as

$$(5) \quad S_K = \sum_{(i,j) \in K} Y_{ij}.$$

Consider now some square K of size $\varphi_0(N) \times \varphi_0(N)$ pixels. Suppose that K contains $S_1^{(K)}(N)$ pixels from all of the particles in the image. The remaining $\varphi_0^2(N) - S_1^{(K)}(N)$ pixels of K contain pure noise. Obviously, $S_1^{(K)}(N) \leq \varphi_0^2(N)$ and, by the definition of K_0 , $S_1^{(K_0)}(N) = 0$. In Langovoy and Habeck (2011b), the following auxiliary statement was proved.

Proposition 1. 1) *If for some K we have $\lim_{N \rightarrow \infty} S_1^{(K)}(N) / \varphi_0^2(N) = 0$, then \bar{Y}_K is a consistent estimate of a .*

2) *In case if $\lim_{N \rightarrow \infty} S_1^{(K)}(N) / \varphi_0^2(N) \neq 0$, it happens that \bar{Y}_K is not a consistent estimate of a .*

This proposition clarifies why a naive approach of estimating a by simple averaging of values over the whole screen leads to a generally inconsistent estimate. The naive estimator is consistent only when a combined size of all particles is negligible compared to the size of the screen. On the other hand, we see that sometimes the averaging procedure is consistent, namely, when the averaging is done over some K such that this square has relatively little intersection with particles. This observation is crucial for our construction of scan estimators.

3 Spatial scan estimators

Definition 1. Let \mathcal{K}_0 be the collection of all $\varphi_0(N) \times \varphi_0(N)$ subsquares of the screen. Set

$$(6) \quad \widehat{K} = \arg \min_{K \subseteq \mathcal{K}_0} \{S_K\}.$$

We define a *spatial scan estimator* (for the lower intensity) as

$$(7) \quad \widehat{a} := \varphi_0(N)^{-2} \sum_{v \in \widehat{K}} Y_v.$$

The spatial scan estimator for noisy patches can be computed via the following algorithm.

- Step 1. Calculate means over all $\varphi_0(N) \times \varphi_0(N)$ subsquares.
- Step 2. Select a neighborhood \widehat{K} with the smallest mean.
- Step 3. Define $\widehat{a} := \varphi_0(N)^{-2} \sum_{v \in \widehat{K}} Y_v$.

Since the initial noisy image can be naturally viewed as a square lattice graph, where k^2 -nearest neighbors correspond to a $k \times k$ subsquare on the screen, we see that the spatial scan estimator (based on a popular method of sliding windows) is a special case of the k -NN scan estimator defined in Langovoy and Habeck (2011b).

Assume that each particle covers a square of size at least $\varphi_1(N) \times \varphi_1(N)$ pixels, and

$$(8) \quad \lim_{N \rightarrow \infty} \frac{\varphi_1(N)}{\log N} = \infty.$$

The above idea can be inverted to get a spatial scan estimator for particle intensities.

Definition 2. Let \mathcal{K}_1 be the collection of all $\varphi_1(N) \times \varphi_1(N)$ subsquares of the screen. Set $\widehat{K} = \arg \max_{K \subseteq \mathcal{K}_1} \{S_K\}$. We define a *spatial scan estimator* (for the higher intensity) as

$$(9) \quad \widehat{b} := \varphi_1(N)^{-2} \sum_{v \in \widehat{K}} Y_v.$$

4 Consistency for bounded noise

In order to establish consistency of our spatial scan estimator, and to derive its rates of convergence, we might impose some additional assumptions on our model. Suppose that for all i and j

$$(10) \quad |\varepsilon_{ij}| \leq M \quad \text{almost surely.}$$

This assumption is realistic in practice. For example, in the cryo-electron microscopy one often has $M = 1$.

The following bound is useful in quantifying behavior of spatial scan estimators (see Langovoy and Habeck (2011b) and Langovoy and Habeck (2011a) for proofs and more details).

Proposition 2. Let K_0 be any square with side $\varphi_0(N)$ completely filled with noise, and let \mathcal{K} be any collection of squares, each of size $\varphi_0(N) \times \varphi_0(N)$. Define $C_1 = 3(b-a)^2$, $C_2 = 12\sigma^2$, $C_3 = 4M \cdot (b-a)$. Then

$$(11) \quad P(\text{some } K \in \mathcal{K} \text{ is chosen over } K_0) \leq \sum_{K \in \mathcal{K}} \exp\left(-\frac{C_1 [S_1^{(K)}(N)]^2}{C_2 |K \setminus K_0| + C_3 S_1^{(K)}(N)}\right).$$

Theorem 1. Suppose that the noise satisfies the above assumptions. If $\lim_{N \rightarrow \infty} \varphi_0(N)(\log N)^{-1/2} = \infty$, the spatial scan estimator (of lower intensities) $\hat{a} = \bar{a}_{\hat{K}}$ is a consistent estimate of a . If $\lim_{N \rightarrow \infty} \varphi_1(N)(\log N)^{-1/2} = \infty$, then the spatial scan estimator (for higher intensities) \hat{b} is consistent for b .

What can be said now about the rate of convergence of \hat{a} to the true a ? In case if K_0 would be known, \bar{a}_{K_0} would in general be an asymptotically efficient estimate of a , and also a $\sqrt{|K_0|}$ -consistent (= $\varphi_0(N)$ -consistent) estimate. One might expect that the rate might be at least $\frac{\sqrt{\log N}}{\varphi_0(N)}$. The additional $\sqrt{\log N}$ -factor is often called a "price for adaptation" in the literature, and estimates that slow down by a logarithmic factor are called "adaptive" (in our case, adaptation is to the extra difficulty of not knowing whether each particular observation belongs to the noise or not). However, it can be shown that the spatial scan estimator actually achieves the parametric rate of convergence. This is a stronger type of adaptivity.

5 Simultaneous detection of multiple particles

The key idea of our particle detection approach is to threshold the noisy picture at a properly chosen level, and then to use percolation theory to do statistics on black and white clusters of the resulting binary image. More details about our approach can be found in Davies et al. (2009), Langovoy and Wittich (2009a), Langovoy and Wittich (2010a), Langovoy and Wittich (2009b), Langovoy and Wittich (2010b) and Langovoy and Wittich (2011). See also related recent work Arias-Castro et al. (2011) and Arias-Castro and Grimmett (2011).

We transform the observed noisy image $\{Y_{i,j}\}_{i,j=1}^N$ in the following way: for all $1 \leq i, j \leq N$,

1. If $Y_{ij} \geq \theta(N)$, set $\bar{Y}_{ij} := 1$ (i.e., in the transformed picture the corresponding pixel is coloured black).
2. If $Y_{ij} < \theta(N)$, set $\bar{Y}_{ij} := 0$ (i.e., in the transformed picture the corresponding pixel is coloured white).

The above transformation is called *thresholding at the level $\theta(N)$* . The resulting array $\{\bar{Y}_{i,j}\}_{i,j=1}^N$ is called a *thresholded picture*.

One can think of pixels from $\{\bar{Y}_{i,j}\}_{i,j=1}^N$ as of colored vertices of a sublattice G_N of some planar lattice. We add coloured edges to G_N in the following way. If any two black vertices are neighbours in the underlying lattice, we connect these two vertices with a black edge. If any two white vertices are neighbours, we connect them with a white edge. The choice of underlying lattice is important: different definitions can lead to testing procedures with different properties, see Langovoy and Wittich (2009a), Davies et al. (2009) and Langovoy and Wittich (2010a). The method becomes especially robust when we work with an $N \times N$ subset of the *triangular* lattice \mathbb{T}^2 . Our goal is to choose a threshold $\theta(N)$ such that $P_0(Y_{ij} \geq \theta(N)) < p_c^{site}$, but $p_c^{site} < P_1(Y_{ij} \geq \theta(N))$, where p_c^{site} is the critical probability for site percolation on \mathbb{T}^2 (see Grimmett (1999)). If this is the case, we will observe a so-called *supercritical* percolation of black clusters within particles, and a *subcritical* percolation of black clusters on the background. There will be a high probability of forming large black clusters in

place of particles, but there will be only little black clusters in place of noise. The difference between the two regions is the main component in our image analysis method.

Algorithm 1 (Detection).

- Step 0. Estimate a and b via the spatial scan estimator. Set $\theta(N) := (\hat{a} + \hat{b})/2$.
- Step 1. Perform $\theta(N)$ -thresholding of the noisy picture $\{Y_{i,j}\}_{i,j=1}^N$.
- Step 2. Until all black clusters are found, run depth-first search on the graph G_N of the $\theta(N)$ -thresholded picture $\{\bar{Y}_{i,j}\}_{i,j=1}^N$.
- Step 3. When a black cluster of size $\varphi_1(N)$ was found, report that there is a particle corresponding to this cluster.
- Step 4. If no black cluster was larger than $\varphi_1(N)$, report that there were no particles.

Suppose that the image contains $\pi(N)$ particles. Assume

$$(12) \quad \lim_{N \rightarrow \infty} \frac{\pi(N)}{\log N} = 0.$$

Let A be a set of pixels corresponding to some particle. Define $\text{thresh}(A) := \{\bar{Y}_{ij} \mid Y_{ij} \in A\}$. For each single particle A , let us say that the particle is *detected* by the Algorithm, if $\text{thresh}(A)$ contains a significant black cluster. For a collection of several particles, the individual significant clusters can possibly merge into bigger clusters. We allow for this in the detection theorem, and leave the task of particle separation for the future pattern recognition work.

Theorem 2. *Let the noise be symmetric and satisfy the above assumptions. Then*

1. *If there were $\pi(N)$ particles in the picture, Algorithm 1 detects all of them with probability that tends to 1 as $N \rightarrow \infty$.*
2. *If there were no particles in the picture, the probability that the Algorithm 1 falsely detects a particle, doesn't exceed $\exp(-C_2(\sigma)\varphi_{im}(N))$ for all $N > N(\sigma)$.*

The constants $C_2 > 0$ and $N(\sigma) \in \mathbb{N}$ depend only on σ and not on particle shapes or exact sizes.

6 Application to particle picking in cryo-electron microscopy

We applied our method to micrographs measured in cryo-electron microscopy. The results are shown in Figure 1. A micrograph of GroEL was downloaded from a public repository for particle picking (http://ami.scripps.edu/redmine/projects/ami/wiki/GroEL_dataset_I). Before running the scan estimators the original 2400×2400 was down-sampled twice to improve the signal to noise and normalized to a maximum intensity of 1. Our scan estimators yield $a = 0.319$ (using a window size of 65) and $b = 0.453$ (using a window with side 9), resulting in the threshold $\theta = 0.386$. After running the percolation analysis on the thresholded image (shown in Figure 1(b)) we delete black clusters occupying less than 30 pixels. The result of the filtering procedure is shown in Figure 1(c). For each of the remaining black clusters, the algorithm claims that there is a particle containing this cluster.

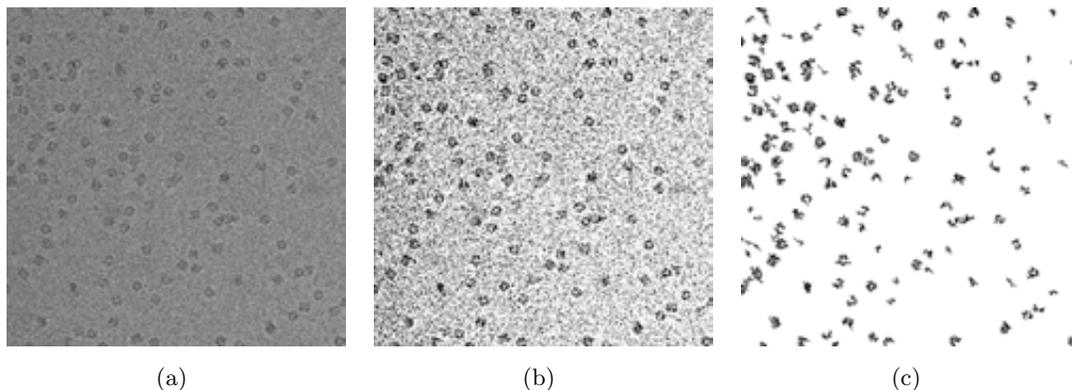


Figure 1: Application to particle picking in cryo-EM. (a): Original micrograph with particles of GroEL. (b): Thresholded image based on spatial scan estimators of the noise and the signal. (c): Thresholded image with random black clusters removed.

References

- E. Arias-Castro and G. Grimmett. Cluster Detection in Networks using Percolation. Submitted, 2011. URL <http://arxiv.org/abs/1104.0338>.*
- E. Arias-Castro, E. Candes, and A. Durand. Detection of an anomalous cluster in a network. Annals of Statistics, 39(1):278–304, 2011.*
- P. L. Davies, M. Langovoy, and O. Wittich. Randomized algorithms for statistical image analysis based on percolation theory. Submitted, 2009.*
- Geoffrey Grimmett. Percolation, volume 321 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, second edition, 1999. ISBN 3-540-64902-6.*
- M. Langovoy and M. Habeck. Adaptive detection of multiple objects in noisy images. In preparation. 2011a.*
- M. Langovoy and M. Habeck. Uniform consistency of spatial scan estimators and k-NN scan estimators. Technical Report. Max Planck Institute for Intelligent Systems, Tübingen, Germany, 2011b.*
- M. Langovoy and O. Wittich. Detection of objects in noisy images and site percolation on square lattices. EURANDOM Report No. 2009-035. EURANDOM, Eindhoven, 2009a.*
- M. Langovoy and O. Wittich. Robust nonparametric detection of objects in noisy images. EURANDOM Report No. 2010-049. EURANDOM, Eindhoven, 2010a.*
- M. Langovoy and O. Wittich. Computationally efficient algorithms for statistical image processing. Implementation in R. EURANDOM Report No. 2010-053. EURANDOM, Eindhoven, 2010b.*
- M. Langovoy and O. Wittich. Multiple testing, uncertainty and realistic pictures. EURANDOM Report No. 2011-004. EURANDOM, Eindhoven, 2011. URL <http://arxiv.org/abs/1102.4820>.*
- M. A. Langovoy and O. Wittich. Detection of objects in noisy images and site percolation on square lattices. Submitted, 2009b.*